Performance studies of scientific applications on High Performance Computing cluster with big-memory footprint

Derrick Lee (UNCC), Garrett Lord (MTU), Jesse Martinez (NMT), Christopher Moore (UK)
Benjamin M. McClelland, Alfred Torrez, Hsing-Bung (HB) Chen, Parks Fields HPC-5
Computer Systems, Cluster & Networking Summer Institute

Introduction

In this project, we installed ScaleMP's vSMP Foundation Software on a IBM x3850 X5 server cluster and setup a virtual SMP machine with a terabyte shared-memory footprint. We have conducted a sequence of benchmark tests that are related to the LANL Data Intensive Super Computing (DISC) project and expect that the vSMP system will be well-suited for large memory problems. Below are the Graph 500 and STREAM Memory benchmark results from the installed system.

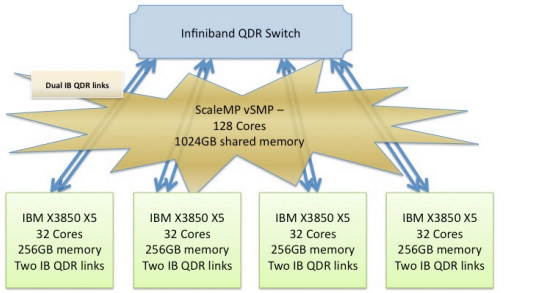

• Once loaded into the memory of each of the system boards, vSMP Foundation aggregates the compute, memory and I/O capabilities of each system and presents a unified virtual system to both the Operating System and the applications running above the OS.

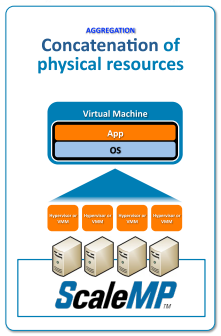

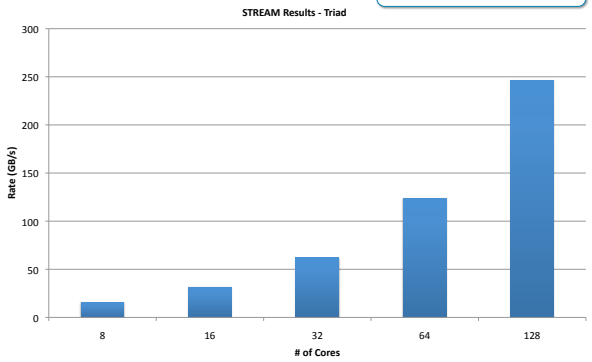

STREAM Memory Benchmark

• (Bottom Left) STREAM benchmark for triad test, performance tends to improve from running 8 threads up to 128 threads.

• vSMP Foundation leverages Intel's Xeon X7550 processor in-board memory bandwidth while maintaining high efficiency for aggregated memory over 4 boards

Graph 500 Benchmark

• (Top Right) Applications of large scale data set graph problems. Interactions are modeled through graph abstractions.

• (Middle Right) The vSMP results are compared to the other Graph500 tests done on other systems. Our current machine has a max of 1,036 billion TEPS which would rank the system 19th against current scale 30 results.

• (Bottom Right) The graph displays the maximum TEPS from the Graph 500 benchmark at various scales sizes and number of cores.

Conclusion

• Preliminary results show that the machine scales linearly and is effective at solving problems with large data sets, such as DNA sequencing and graph converting problems.

• Current results show how promising the vSMP Foundation is for increasing scale sizes due to high memory bandwidth.

•As with vSMP Foundation, the memory location is transient, and as one places threads on a separate core, each thread can leverage (and benefit) from the maximum memory bandwidth on that core/socket.

• In the future, we plan to apply our testing experience and collective knowledge to work with LANL's Biology, Astrophysics, Chemistry, Cyber-security, and Visualization communities on their Data Intensive Computing projects.

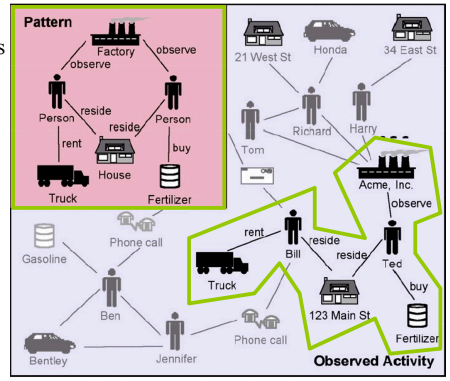

Graph-based technologies for intelligence analysis, T. Coffman, S. Greenblatt, S. Marcus, Commun. ACM, 47(3):45-47,2004.

• This graph represents a real life situation that simulates multiple vertices. Linking and extracting relevant data in a timely manner can be addressed with large memory footprint systems.

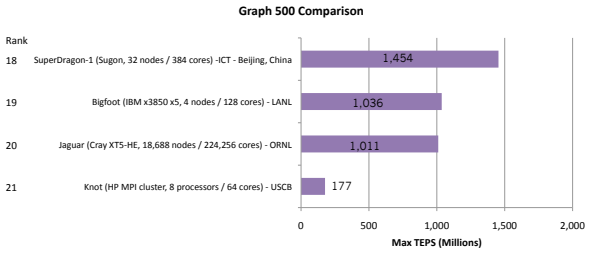

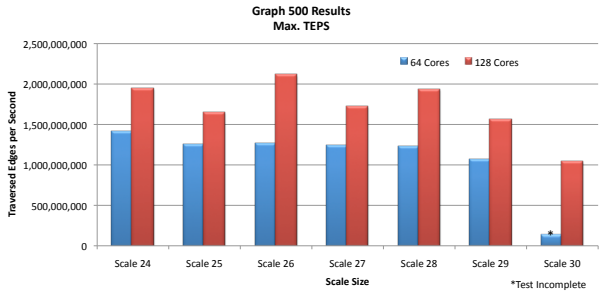

Contact Information:
HB Chen (hbchen@lanl.gov)

LA-UR-11-11209